 MOTION IMAGERY STANDARDS BOARD STANDARD Compression for Infrared Motion Imagery	MISB ST 0404.1 27 February 2014
--	--

1 Scope

Infrared imaging systems are a common sensor of choice for many imaging platforms, including unmanned autonomous systems (UAS). They are capable of providing high contrast imagery in the day and night and maintain a dynamic range that exceeds that which is available with broadcast format visible imaging systems, such as high definition 720p and 1080p. While the broadcast market has driven the requirements and advancement for compressed video streaming systems, scarce attention has been paid to the development of standard hardware interfaces and compression algorithms that are suited for optimum use with infrared imagery.

This Standard addresses unique issues that pertain to optimizing the quality and utility of data that are acquired with infrared imaging systems. While infrared imaging systems routinely provide 14-bit dynamic range pixels, standard broadcast format interfaces that are leveraged on many UAS systems are limited to eight or ten bits. State of the art compression algorithms, including JPEG 2000 and H.264, now support monochrome (4:0:0) compression at 14-bits. This advancement allows for the retention of dynamic range information that is traditionally lost when utilizing eight or ten bit streaming systems.

This Standard provides recommendations and required modifications to existing systems and software to leverage 14-bit compression capabilities. These modifications include the processing and display of 14-bit imagery. Recommendations are also made for formatting and interfacing infrared imagery with existing broadcast format eight and ten bit systems. Given certain types of system constraints, preferred methods are provided for configuring and performing compression that best preserves data integrity after high bit-depth conversion.

2 References

2.1 Normative References

The following references and the references contained therein are normative.

- [1] ISO/IEC 13818-2:2000, Information technology - Generic coding of moving pictures and associated audio information: Video
- [2] ITU-T Rec. H.264 (04/2013), Advanced Video Coding for Generic Audiovisual Services

2.2 Informative References

- [3] ISO/IEC 15444-1:2004, Information Technology – JPEG 2000 Image Coding System: Core Coding System
- [4] SMPTE ST 292-1:2012, 1.5 Gbit/s Signal/Data Serial Interface
- [5] MISB RP 0904.1 H.264 Bandwidth/Latency/Quality Tradeoffs, Feb 2014
- [6] MISB RP 0401 Infrared Motion Imagery System Matrix
- [7] MISB RP 0402.5 Parallel Interface for Infrared Motion Imagery, Feb 2014
- [8] MISB ST 0403.2 Digital Representation and Source Interface formats for Infrared Motion Imagery mapped into a 1280 x 720 format Bit-Serial Digital Interface, Feb 2014
- [9] N. J. McCaffrey, F. S. Pantuso, “Very Low Cost Histogram Based Contrast Enhancer Utilizing Fixed Point DSP Processing”, Proc. SPIE 3303, PP 36-43.
- [10] V.E. Vickers, “Plateau Equalization Algorithm for Real –Time Display of High Quality Infrared Imagery”, Opt. Eng. 35(7), pp 1921-1926, 1996

3 Acronyms

IR	Infrared
NIR	Near-Infrared
NTSC	National Television Systems Committee
SWIR	Short Wave Infrared
MWIR	Mid-Wave Infrared
LWIR	Long-Wave Infrared
HD	High Definition
FRExt	Fidelity Range Extensions
MTF	Modulation Transfer Function
ADR	Airy Disc Radius

4 Revision History

Revision	Date	Summary of Changes
ST 0404.1	02/27/2014	<ul style="list-style-type: none"> • Updated references; EARS format

5 Introduction (Informative)

For a long period, infrared imaging systems have played an integral role. This is partly because of the day/night capabilities of certain infrared bands as well as the target discrimination capabilities that can be achieved with the infrared bands. The potential benefits of infrared imaging have been somewhat constrained however, as the current system infrastructure that infrared cameras interface with has been designed around visible imaging systems that have slightly different characteristics and features. The dynamic range and bit depths used in visible broadcast format systems have acted as a constraint on infrared systems. This constraint has been partially lifted with the availability of modern compression algorithms that support 14-bit monochrome imagery, such as JPEG 2000 and H.264 w/ Fidelity Range Extensions. With the

advent of this recent compression algorithm enhancement that allows for 14-bit monochrome operation, the native capabilities of infrared sensors can be more fully realized.

Traditionally, infrared cameras have, internal to the camera, converted its wide dynamic range sensor information to a scaled version more compliant with either 8 or 10-bit visible broadcast video hardware and software. This Standard outlines options that are available for the system designer in the configuration of an infrared streaming system. This includes systems that leverage legacy MPEG-2 compression (see Figure 1), as well as JPEG 2000 and H.264 based compression, that support the full native dynamic range capabilities of infrared cameras (see Figure 2). An issue that can be noted from the two figures is the location of 14-bit monochrome to 24-bit color scaling. In legacy systems, this function has traditionally been performed inside the infrared camera. When leveraging newer 14-bit compression capabilities, the process of converting from 14-bit monochrome to 24-bit color will need to be added to viewing stations and software on the backend. This may have significant implementation ramifications.

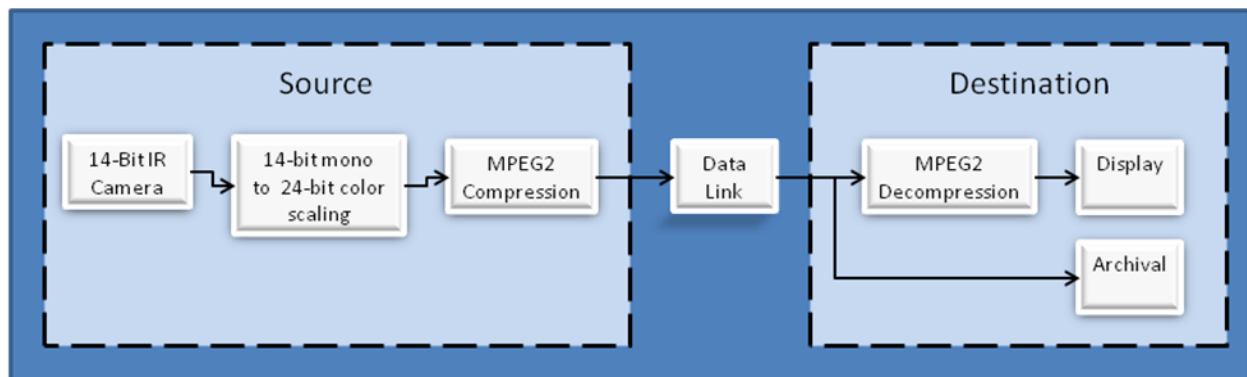


Figure 1: Legacy Infrared System Configuration

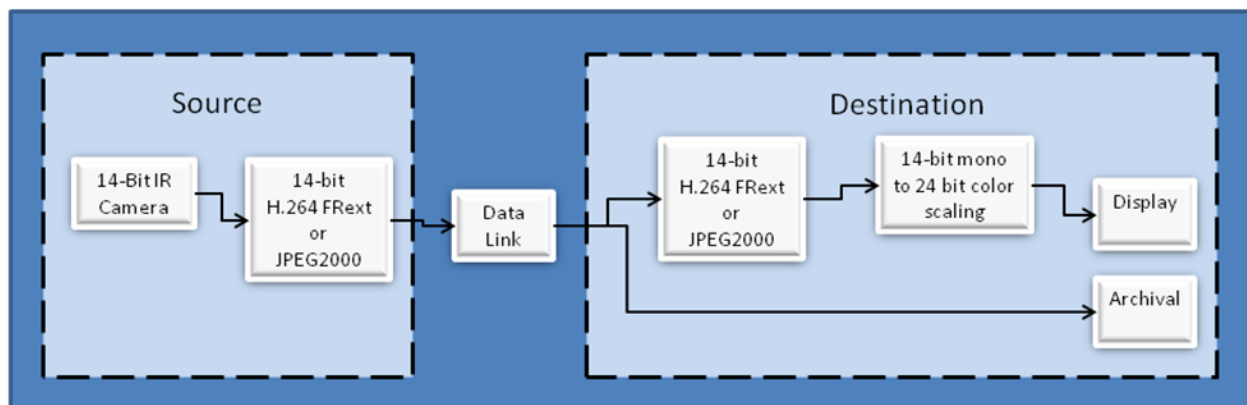


Figure 2: Infrared system leveraging 14-bit compression.

Background information on the physics of infrared imaging is presented to show how the configuration of a system and the operational environment of a system can have an effect on the

compression performance that is achieved. Finally, differences between the different infrared bands are discussed, also to show how specific band characteristics relate to compression effectiveness that is achieved across the different infrared bands.

6 Infrared Imaging (Informative)

On the surface, infrared imagery may appear to be equivalent to monochrome visible imagery with just a higher dynamic range. There are subtle differences; however, that if well understood can lead to a more efficient and useful data stream when compression is required. Issues such as optical diffraction, pixel size, and atmospheric turbulence have a profound effect on the modulation transfer function (MTF) of an imaging system. As a result these parameters have a direct impact on the “compressibility” of the imagery. Optical diffraction and atmospheric turbulence are both functions of wavelength and therefore have a different effect in the visible wavelength band than in the different infrared wavebands.

6.1 *Description of infrared*

The infrared region, as it relates to this Standard, is typically defined as electromagnetic radiation ranging from just below 1 micron to approximately 14 microns. This region of the electromagnetic spectrum includes the near or short-wave, mid-wave, and long-wave bands. The region above approximately 2.4 microns in wavelength is typically referred to as the thermal bands because items that are close to ambient temperature emit blackbody radiation in these bands. As a result, people and targets tend to show up well in the thermal bands. This is what gives mid and long wave cameras day and night time imaging capabilities. They don’t require solar radiation for illumination. In the short wave infrared, the majority of signal that is available is generated from solar radiation. Other items of interest that are common in the short wave include laser radiation, such as that from laser designators and laser weapon systems.

6.2 *Dynamic Range*

A characteristic feature of infrared imaging systems is their typically high dynamic range. Many thermal infrared scenes, however, are typically characterized by a low scene dynamic range. This is caused by the fact that in many natural scene settings, the temperature of items in the scene is typically close to uniform. While a typical mid-wave sensor may have a dynamic range of 300 to 400 degrees C, everything in a natural outdoor scene may very well be within ten or a few tens of degrees C. As a result, a large portion of the dynamic range can go unused.

Figure 3 shows an example of this occurrence where the histogram on the right shows almost the entire scene residing in a small sliver of the 14-bit dynamic range space. Figure 4 shows the same scene after proper scaling.

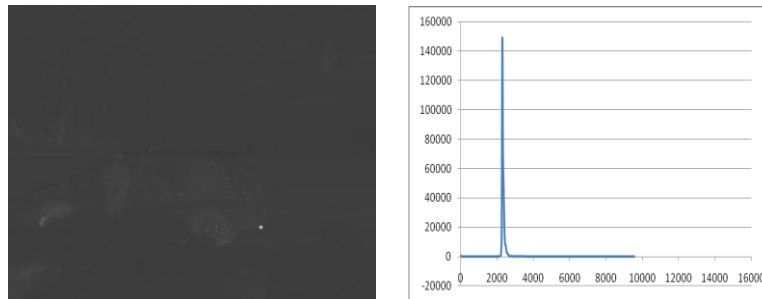


Figure 3: 14-bit IR image linearly truncated to 8-bits, with histogram

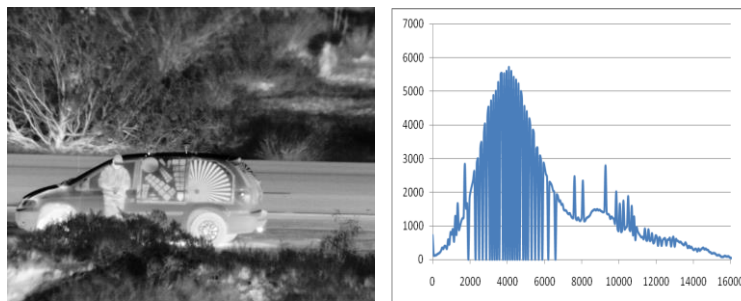


Figure 4: Mid-wave infrared image that is properly scaled

Sensor dynamic range defines the ability to capture both low and high light levels at the same time. The majority of commercial visible systems capture either 8 or 10 bit monochrome or color with some scientific sensors achieving 12 or more bits of dynamic range. Infrared sensors typically provide 14-bits of sensor dynamic range. Scene dynamic range varies depending on the amounts of both reflected and emitted energy and the scene content. For scenes with mostly reflected energy, such as visible and near-IR, scenes typically exhibit very high dynamic range (in the daytime). For scenes with mainly emitted energy (thermal), natural scenes are typically lower in scene dynamic range. This can change however, when scene content such as hot engines, explosions, muzzle blasts, etc. are present.

In Section 7, methods for converting the 14-bit image to lower bit-depth monochrome or color images are presented. This process is required when leveraging legacy MPEG-2 compression systems, and also when displaying the decompressed results of the 14-bit images when leveraging JPEG 2000 and H.264. The conversion for display purposes is required as displays typically operate in the 24-bit color space.

This wide dynamic range allows for the potential of capturing detail that would otherwise be lost in a lower dynamic range acquisition system. When using standard 8-bit compression systems, such as MPEG-2, details that are potentially desirable can be lost in the compression process. As the scaling process is required to be performed before the compression is carried out, the portions of the dynamic range that contain objects of interest must be expanded properly. If the scaling is not performed correctly, desired information can be lost and is no longer retrievable within the 8-bit image. In the case of 14-bit compression, an end-user viewing the imagery has the opportunity to rescale the imagery to enhance the areas of the dynamic range that are of interest.

For instance, for a scene with both a blast explosion and shadowed areas, both the blast details and the shadows can be enhanced, but usually not at the same time. This is one of the key benefits and justifications for implementing 14-bit IR compression.

Figure 5 shows an example where a 14-bit image with both hot and cold components has been scaled using a different technique. In the first case, the scaling has enhanced the detail of the torch in the center of the frame as well as the detail on the soda bottle and coffee pot, which has ice water in it. In the second scene, the scaling has enhanced the detail on most of the objects that are at room temperature. Cables and other items that are not readily obvious in the first scene now stand out. The water with ice that has been shot across the screen shows up much more completely. In addition, detail on the torch, the soda bottle and the coffee pot with ice water has been lost. If an end user has access to the 14-bit version of this image, either scaling can be achieved depending on the portion of the scene that is of interest. If only the 8-bit version is available, the information of interest may or may not be lost depending on the choice of scaling.

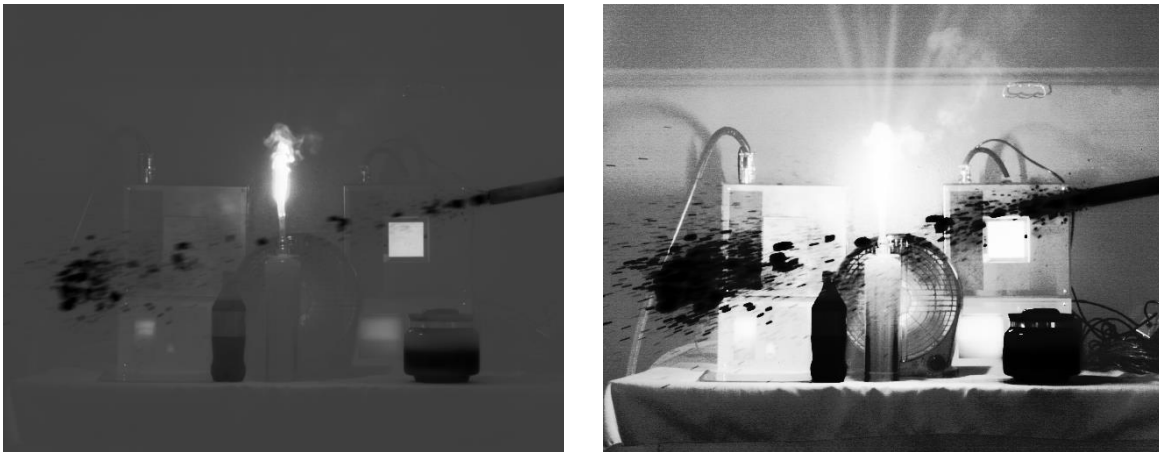


Figure 5: Comparison of different scaling results

7 Infrared Systems (Informative)

To achieve optimum results when compressing infrared imagery, it is important to understand and factor in all of the key issues that go into the formation of an infrared image. The following sections describe the different infrared bands, camera interface issues, resolution definitions, compression options, display algorithms, and mechanisms for manipulating the infrared imagery when dealing with highly constrained bit-rates.

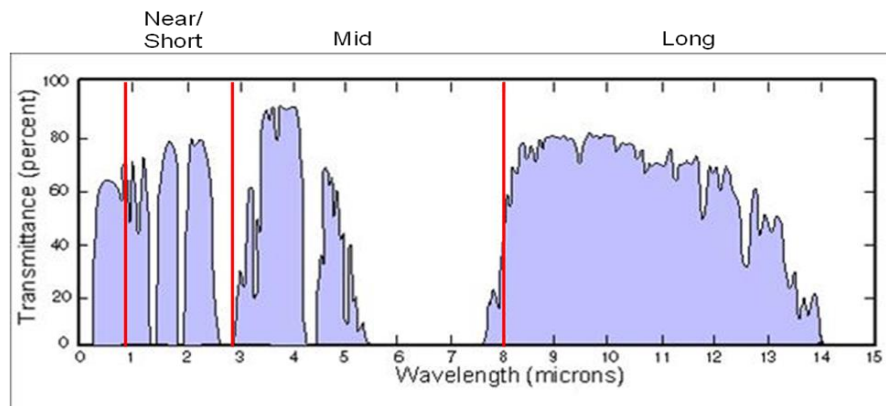
7.1 Band Definitions

Infrared imaging typically spans multiple sub-bands, as previously described. There are varying definitions for these bands depending on whether the focus is on atmospheric transmission, sensor characteristics, etc. Table 1 shows a common description for the overall bands and for the portion that is practical for imaging, due to absorption, sensor material sensitivity, etc.

Table 1: Infrared band descriptions

Band	Wavelengths	Typical Imaging Wavelengths
Near/Short Wave IR	750 – 3000 nm	900 – 1700 nm
Mid-Wave IR	3000 – 8000 nm	3000 – 5000 nm
Long Wave IR	8000 – 14000 nm	8000 – 12000 nm

Figure 6 shows the atmospheric transmission curves for the bands of Table 1. As can be seen, there are significant absorption areas, due to water vapor, carbon dioxide and other gases. In some of these regions, almost no transmission occurs, such as the area from 5 to 8 microns.

**Figure 6: Atmospheric transmission**

7.2 Hardware interfaces

Interfaces for infrared camera systems are defined in the MISB RP 0402 [7] and MISB ST 0403 [8]. ST 0403 outlines the use of the SMPTE ST 292 [4] serial digital interface for use with infrared systems. The ST 292 interface defines a transmission mechanism for 10-bit imagery. To interface 14-bit infrared imagery through this 10-bit interface requires manipulation of the infrared pixels both before and after the interface. This is usually achieved by placing some of the bits from each pixel into the color difference space. ST 0403 defines a mechanism for this type of mapping. Of concern here is when implementing compression systems that interface with the ST 292 interface outlined in ST 0403, the infrared imagery must be reconstructed into its proper intended format prior to compression.

7.3 Compression Issues

The goal of compression is to retain as much desired information content as possible while throwing away redundancy and information that is less relevant. The following sections outline issues relating to the compression of infrared imagery and provide insight into how systems can be configured to optimize infrared compression results for particular situations.

7.3.1 Resolution and Compressibility of Infrared Imagery

The compressibility of imagery is a function of the information content or entropy that is contained in an image sequence. A number of factors go into determining how much information content is contained within a specific sequence. Major factors include scene content and motion, pixel count, frame rate, optical resolution, atmospheric turbulence, pixel size, pixel depth, sensor non-uniformity, and noise. While scene content and image motion, pixel count and frame rate are typically the focus of attention, the other parameters can also have a significant effect on the compressibility of a video stream. Figure 7 shows the major components involved with the formation of an image.

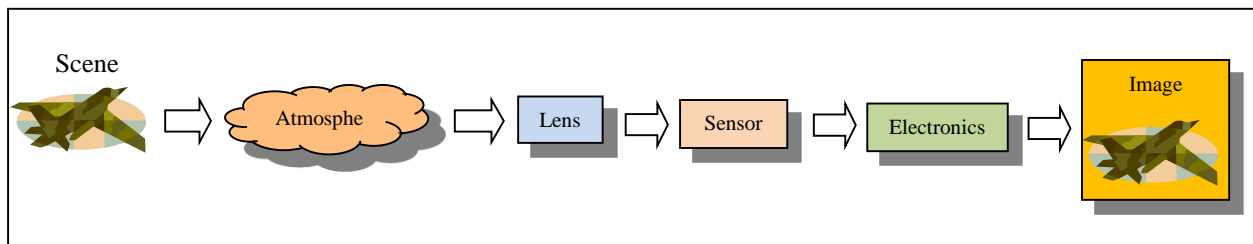


Figure 7: Image Formation Chain

When a scene is captured, the content and environmental effects on the content and sensor affects the frequencies and contrast within an image. If a scene contains a significant amount of small detail, such as leaves, blades of grass, crowds of people, etc. i.e. high entropy, it will be more difficult to compress. Conversely, when the scene has regions of smooth continuous areas, such as clear sky i.e. low entropy, it will be easier to compress. In addition, atmospheric turbulence is a stochastic process that tends to degrade and reduce image resolution, which at the intra-frame level is more efficient to compress. Over time however, turbulence causes image distortion and introduces motion, which can make inter-frame or predictive frame calculations more difficult and compression less efficient. An aspect of atmospheric turbulence that comes into play with infrared compression is that the degradation scales with wavelength. As a result, while turbulence affects imaging in the visible band to a significant extent, it does not affect the mid-wave band very much and even less in the long wave band.

Atmospheric turbulence also has a geometric property that is important to understand, and impacts imaging depending on distance from the ground. As the strongest turbulence is usually found near the ground, approximately 30 feet or so, imaging across the ground will usually observe degradation due to turbulence. Imaging further above the ground will usually be subject to less degradation, where finally satellite or airborne platform imaging looking down at the ground see the least. This phenomenon results from light bent by strong turbulence statistically will be unlikely to hit the relatively small aperture of an optical collection platform that is far away. While these generalities may vary somewhat, they are good rules of thumb to follow.

7.3.2 Image Sampling

The selection of scene and the state of the atmosphere is rarely under the control of a system designer; however, knowledge of the optics, sensor, and processing electronics can help predict expected performance and determine appropriate system operating points. With respect to a

camera lens, the Airy disc radius (ADR) is a common metric for defining the limiting optical resolution. The Airy disc can be calculated as follows:

$$\text{ADR} = \frac{1.22 \cdot \lambda \cdot f}{D} = 1.22 \cdot \lambda \cdot f/\# \quad \text{Eq. 1}$$

Where λ is the wavelength, f is the focal length of the optic and D is the aperture diameter of the optic. The f -number of a lens ($f/\#$) is equivalent to the focal length divided by the aperture diameter. The ADR can be thought of as the smallest feature size that can be passed through the optic. The relationship of the ADR to the pixel size is a key parameter with respect to compressibility of an image. The “pixel” is essentially a sampling tool and the ADR defines the maximum detail that can be sampled with the pixel.

Note that the ADR increases in size linearly with wavelength. Long-wave infrared cameras, with a similar f -number, will have a larger ADR than a mid-wave camera, which will have a larger ADR than a short-wave camera. As such, assuming a similar pixel size, one would expect the long-wave IR to compress better than the mid-wave, which would compress better than the short-wave. Table 2 shows the pixel sampling relationship with the optical blur size for IR cameras with 25 μm pixels and an $f/4$ optic, which are reasonably common.

Table 2: Comparison of optical blur and pixel sampling

IR Band	ADR	Pixel Size	Sampling Ratio
Short-wave	6.3 μm	25 μm	0.25
Mid-wave	20 μm	25 μm	0.8
Long-wave	45 μm	25 μm	1.8

As mid and long wave cameras almost always have a fixed f -number (via cold stop), there is usually no mechanism for changing the relationship between the optical blur size and the pixel size. For a short-wave IR camera, however, there frequently exists the opportunity to change the f -number through the iris setting on the lens. The iris therefore can act as a tuning parameter for image compression i.e. fine detail can be compromised if high compression ratios are required. The iris essentially acts as a lowpass filter on the light passing through the lens. The caveat here is that enough light for generating a good exposure must be passed through the optic.

7.3.3 Sensor Noise and Non-Uniformity

Like all image sensors, infrared sensors generate noise and exhibit a certain level of non-uniformity across the sensor. Both noise and non-uniformity are frequently manifested as single pixels with values that exceed an expected difference from their neighboring pixels. The expected difference is usually defined by the modulation transfer function of the optical system that is projecting light onto the sensor. Most infrared sensors have a certain level of pixel non-uniformity due to imperfect manufacturing issues. It is usually minimized during operation through a calibration procedure that determines an offset and gain correction for each pixel in the sensor array. As the calibration values can vary over time with environmental and other

changes, operating procedures should dictate proper implementation of calibration procedures so as to minimize the detrimental effects of non-uniformity errors.

When noise or non-uniformity is present, the high frequency content in the image is increased, which is more difficult to compress. One mechanism for reducing these unwanted image artifacts is by pre-filtering the image. Care must be taken to filter out only the unwanted features in the image, while minimizing the reduction of desirable high frequency content.

8 Infrared Compression Options

This Standard identifies allowed compression algorithms for infrared compression. For compatibility with legacy systems MPEG-2 [1] compression is allowed for 8-bit 4:2:0 color applications. For new systems, H.264/AVC [2] is the preferred compression type. The *Advanced Video Coding Fidelity Range Extensions* in H.264, called High Profile, is recommended over MPEG-2 for providing higher bit depth, monochrome operation, and superior compression performance. The High 4:4:4 Profile operated in monochrome mode is recommended for 14-bit depth magnitude monochrome operation, and the H.264 compression performance.

For applications where H.264 is not available for 8 through 14-bit monochrome compression, JPEG 2000 is allowed as an acceptable alternative. Figure 8 provides a diagram of possible system implementations for each type of potential stream output format. For all applications, higher bit depths are preferred when available bandwidth allows for their implementation. For low bit-rate applications there may be scenarios where 8-bit monochrome compression is preferred because of bandwidth constraints.

8.1 H.264/AVC

Requirement	
ST 0404.1-01	For new infrared compression systems H.264/AVC [2] shall be utilized for 8 -14 bit monochrome (4:0:0) and 24-42 bit color (4:2:0, 4:2:2, and 4:4:4) applications.

8.2 MPEG-2

Requirement	
ST 0404.1-02	When compressing infrared motion imagery with MPEG-2 [1] the Main Profile Main Level (MP@ML) shall be used for 720x480/30Hz and 720x576/25Hz imagery.
ST 0404.1-03	When compressing infrared motion imagery with MPEG-2 [1] the Main Profile High Level (MP@HL) shall be used for 1280x720/60Hz imagery.

8.3 JPEG 2000

For cases where H.264 High Profile is not available for use on a system with 8 through 14 bit monochrome infrared imagery, JPEG 2000 [3] is an acceptable alternative for use with this

imagery. JPEG 2000 provides comparable performance to H.264 in the High 4:0:0 Intra-frame only mode.

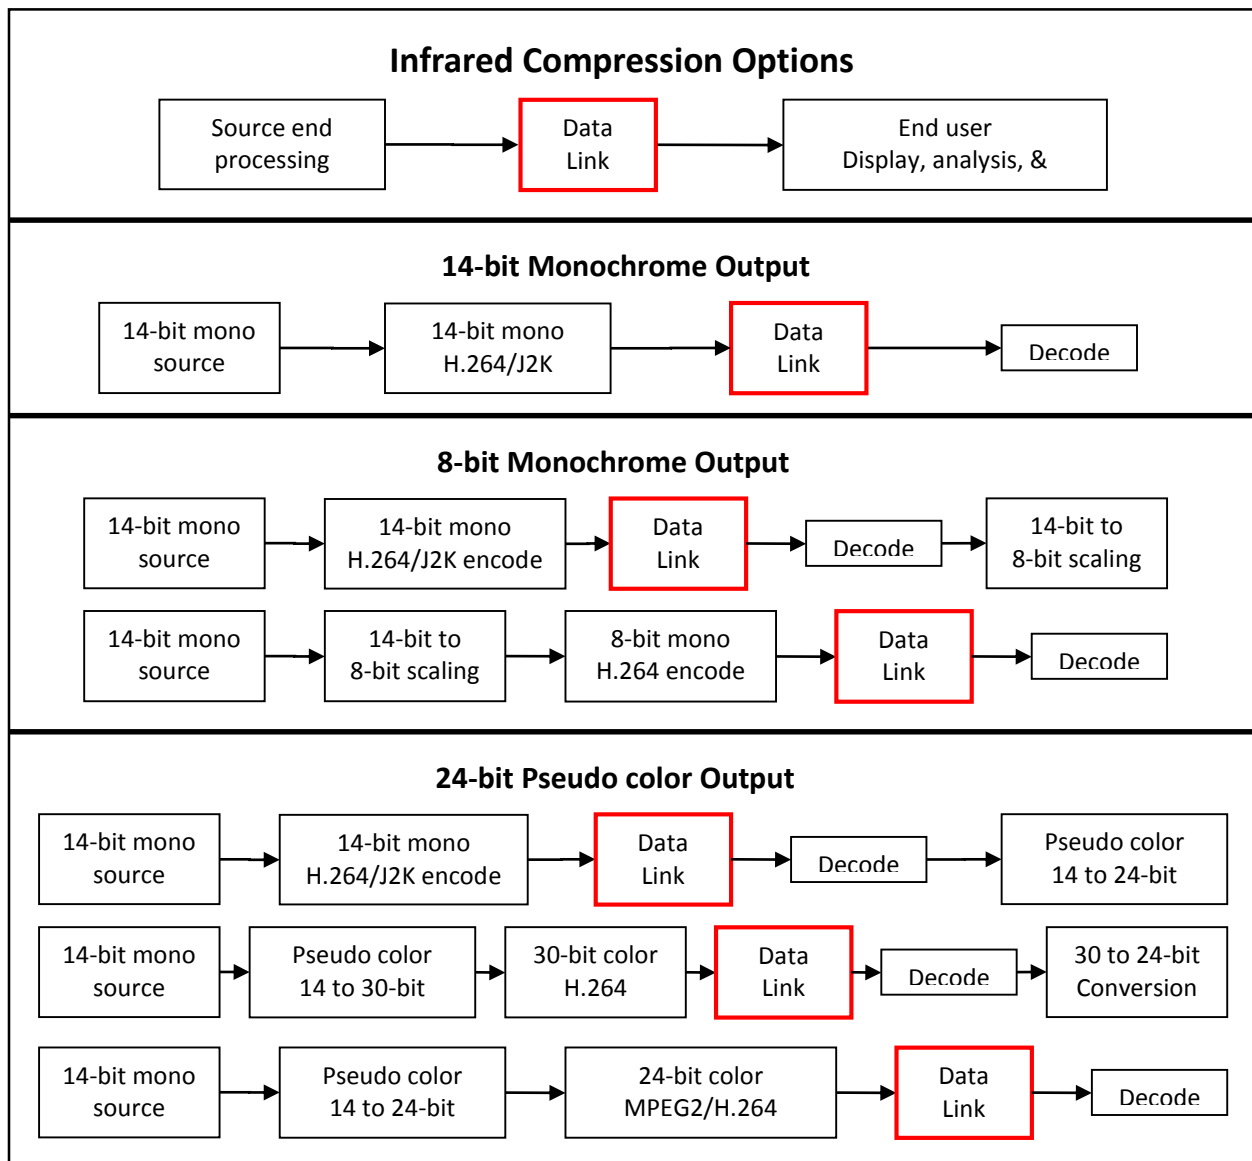


Figure 8: Infrared compression options based on desired output format

9 Bandwidth/Latency/Quality Tradeoffs (informative)

MISB RP 0904 [5] discusses methods for delivering high definition motion imagery over a limited bandwidth channel by changing the image spatial format (scale/crop) and the temporal frame rate. In situations where a low bit-rate stream is required, compression of full resolution 14-bit infrared motion imagery may not produce desired results. When compression rates reach approximately 40:1 or more for intra-frame compression, and 120:1 for inter-frame compression

certain tradeoffs may warrant different solutions. Table 3 lists options for reducing the information content of the imagery prior to compression.

Table 3: Options for pre-processing imagery for low bit rate applications

Method of Bit-rate Reduction	Image Quality Effect
Binning	Lowers pixel resolution by a factor of four for each level of binning. Causes an increase in neighboring pixel contrast resulting in a less than factor-of-four improvement in bit-rate reduction.
Windowing (Cropping)	Narrows the overall field of view while keeping the same pixel field of view. Acceptable if desired scene content remains within field of view.
Low Pass Filtering	Lowers high frequency content and softens edges. Can be effective in instances with poor non-uniformity correction.
Conversion to 8-bit	Cuts the byte count for the image in half. Can be very effective for images with low scene dynamic range. For scenes with high dynamic range, conversion may result in loss of desired detail.
Frame Rate Reduction	Throws away temporal information. Effective for scenes with content that has low temporal bandwidth. For scenes with high temporal scene motion, lower frame rates can cause challenges and inefficiencies for inter-frame compression.

9.1 *Bit Depth Conversion and Display of Infrared Imagery (Informative)*

Whether converting 14-bit infrared imagery to eight bits for input to a compression system, or for display, there are a variety of approaches for performing the conversion. The following sections introduce common methods used to perform this type of conversion. There are two main categories of conversions: contrast stretching and pseudo-coloring.

9.1.1 Contrast Stretching

Converting a 14-bit image to an 8-bit monochrome image causes, in most cases, the loss of some information. Given that some loss is inevitable, the goal is for a conversion that maximizes the subjective quality of the resulting image. Many common approaches for bit-depth conversion are based on the histogram of an image. Examples include plateau equalization, histogram projection, etc. Many of these algorithms are well understood and leveraged in most fielded systems in one manner or another. The informative references list two of these types of algorithms and capabilities [9][10].

9.1.2 Conversion to Color

While infrared cameras generally produce monochrome imagery, the majority of compression codecs and display systems require a color input. As a result, for typical system implementation, a mechanism for converting the monochrome source imagery into a color format is required.

Look-up tables are a quick and efficient method for performing this conversion. In addition, a myriad of monochrome to color palettes are available to provide different effects for enhancing different aspects of a scene. Figure 9 shows an example of different palettes that have been applied to the same image.



Figure 9: Examples of pseudo-color palettes for infrared imagery

9.1.3 Group of Pictures (GOP) Coding Structure

A key “tuning” parameter that can be used to vary the bandwidth, latency, and quality of a compressed infrared data stream is the Group of Pictures (GOP) coding structure. For low latency applications, an intra-frame (I-frame) only configuration is ideal; however, the tradeoff is compression efficiency. Applications that require low bandwidth more efficient compression would likely choose using predictive coding structures i.e. a mix of predictive (P) and bi-directional (B) frames. The tradeoffs in using P and B frames is increased latency and “brittleness” with respect to fast, high dynamic transient scene changes. Table 4 provides recommendations for data rate for various 14-bit infrared resolution and frame rate configurations. The two GOP coding structure implementations include I-frame only and a 16 frame GOP IBBP structure. These are considered appropriate for many real-time implementations.

Table 4: Data Rates for I-frame and IBBP GOP coding structures

	Width (pixels)	Height (pixels)	Frame Rate (FPS)	Uncompressed Bit Rates (Mb/sec)	Target Tracking Bit-rates (Mb/sec)		Critical Imaging Bit-rates (Mb/sec)	
					I-Frame	16 Frame IBBP GOP	I-Frame	16 Frame IBBP GOP
IR-MISM Levels 1-3								
Maximum	180	144	60	21.773	0.544	0.181	1.089	0.363
Minimum	160	120	5	1.344	0.034	0.011	0.067	0.022
Nominal	160	120	30	8.064	0.202	0.067	0.403	0.134
IR-MISM Levels 4-6								
Maximum	360	288	60	87.091	2.177	0.726	4.355	1.452
Minimum	320	240	5	5.376	0.134	0.045	0.269	0.090

ST 0404.1 Compression for Infrared Motion Imagery

Nominal	320	240	30	32.256	0.806	0.269	1.613	0.538
IR-MISM Levels 7-9								
Maximum	720	576	60	348.365	8.709	2.903	17.418	5.806
Minimum	640	480	5	21.504	0.538	0.179	1.075	0.358
Nominal	640	480	30	129.024	3.226	1.075	6.451	2.150
IR-MISM Levels 10-12								
Maximum	1280	1024	60	1101.005	27.525	9.175	55.050	18.350
Minimum	1024	720	5	51.610	1.290	0.430	2.580	0.860
Nominal	1024	1024	30	440.402	11.010	3.670	22.020	7.340
IR-MISM Levels 13-15								
Maximum	2048	1152	60	1981.809	49.545	16.515	99.090	33.030
Minimum	1920	1080	5	145.152	3.629	1.210	7.258	2.419
Nominal	1920	1080	30	870.912	21.773	7.258	43.546	14.515
IR-MISM Levels 16-18								
Maximum	3840	2160	60	6967.296	174.182	58.061	348.365	116.122
Minimum	2048	1152	5	165.151	4.129	1.376	8.258	2.753
Nominal	2048	2048	5	293.601	7.340	2.447	14.680	4.893

Notes: The above resolutions represent the majority of commercial IR sensor resolution formats. The 1024x720 resolution is derived from cropping a 1024x1024 IR sensor as required for formatting to a 1280x720 interface as defined in MISB ST 0403 [8].

10 Infrared Compression Options

Table 5 lists different options for compressing infrared imagery. Level D compression represents the typical legacy system that leverages MPEG-2 compression. These systems typically ingest analog NTSC video, but may also ingest video from other formats. Level D represents the poorest image quality solution of the options that are listed. Level A compression, on the other hand, represents the solution that holds the potential for obtaining the highest level of image quality. It is the only mode that operates as a full 14-bit solution from sensor output through compression to display. Formatting of infrared imagery for the SMPTE ST 292 interface is described in MISB ST 0403.

Table 5: Infrared compression options

IR Compression Options	Compliance Description	Input Color Format	Implementation
Level A	Fully Compliant	14-bit 4:0:0	JPEG 2000 or H.264 FExt Profile IDC Level 244, High 4:4:4
Level B	Partially Compliant	10-bit 4:2:2	Scaled and converted to 10-bit 4:2:2 color and compressed with H.264
Level C	Less Compliant	8-bit 4:2:0	Scaled and converted to 8-bit 4:2:0 color and compressed with H.264
Level D	Minimally Compliant	8-bit 4:2:0	Scaled and converted to 8-bit 4:2:0 color and compressed with MPEG-2

Table 4 provides general guidance on compression stream settings for infrared encoding systems that are leveraging Level A infrared compression as defined in Table 5. For any constant bit rate setting, the achieved quality of compressed video will vary depending on a variety of factors that include scene content, atmospheric turbulence, wavelength of imaging, optics configuration, quality of non-uniformity correction, pixel and lens matching, etc. Given these and other factors, the values defined in Table 4 provide good general guidance for constant bit-rate settings given the performance of current H.264 codecs at 14-bits. For applications where environmental factors dictate, it may be appropriate to adjust these settings either up or down.

This Standard defines requirements for the compression of infrared motion imagery. H.264 FExt and JPEG 2000 are both capable and approved for use in compressing 14-bit infrared motion imagery. Performance of the two in intra-frame mode are comparable in bit rate; however they display different characteristic artifacts when compression ratios become significant. H.264 in inter-frame compression mode provides higher compression ratios than single frame JPEG 2000 compression. As of the release of this version of the Standard, inter-frame compression for JPEG 2000 has not shown the level of maturity required to warrant inclusion at this time.

11 Appendix (Informative)

The following 'C' code is a sample general purpose algorithm for scaling 14-bit monochrome infrared imagery into an 8-bit monochrome output image. The algorithm is a modified version of the histogram equalization algorithm in which the calculated weighting terms are scaled with a square root function.

```
void Sample14bitScaling(unsigned short* ImageIn, long Rows, long Columns, unsigned char* ImageOut)
{
    int i, Equalize[16384];
    double Histo[16384], sum;

    for(i=0; i<16384; i++)                /*< Zero out histogram array */
    {
        Equalize[i] = 0;
        Histo[i] = 0.0;
    }

    for(i=0; i< Rows*Columns; i++)        /*< Calculate histogram of image */
    {                                       /*< and store result in Equalize array */
        Equalize[16383&ImageIn[i]]++;
    }

    sum = 0.0;
    for(i=0; i<16384; i++)
    {
        Histo[i] = sqrt((double)Equalize[i]);    /*< Scale equalization terms via square root*/
        sum += Histo[i];
    }

    for(i=1; i<16384; i++)
    {
        Histo[i] += Histo[i-1];
    }

    for(i=0; i<16384; i++)                /*< Normalize summed values to output gray level space */
    {
        Equalize[i] = (int)(( 255.0*Histo[i])/(double)sum);
    }

    for(i=0; i< Rows*Columns; i++)        /*< Remap the input pixels to create the equalized image */
    {
        ImageOut[i] = (unsigned char) (Equalize[16383&ImageIn[i]]);
    }
}
```